DOCUMENT RESUME

ED 328 215 IR 014 809

AUTHOR Elsom-Cook, Mark
TITLE Guided Discovery Tutoring and Bounded User Modelling. CITE Report No. 13.
INSTITUTION Open Univ., Walton, Bletchley, Bucks (England). Inst. of Educational Technology.
PUB DATE 88
NOTE 18p.; A product of the Centre for Information Technology in Education (CITE). For related reports, see IR 014 806-808 and IR 014 811-813.
PUB TYPE Reports - Evaluative/Feasibility (142)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Computer Assisted Instruction; *Discovery Learning; Expert Systems; Instructional Systems; Interaction; Models; *Programed Tutoring; Teacher Student Relationship; *Teaching Methods

ABSTRACT

This paper discusses a range of teaching styles possible for use in Intelligent Computer Aided Instruction (ICAI) and suggests that they may usefully be thought of as points within a framework of guided discovery teaching methods. The nature of guidance in tutoring and possible ways of providing guidance are discussed in the context of teacher-student relationships and the teaching interaction that occurs during an educational exchange. The use of Bounded User modelling within a guided discovery framework is introduced, and an application of the technique in a Lisp teaching system is described. Also mentioned are various instructional programmed tutoring systems including GUIDON, WHY, WEST, and IMPART. Extensions to the system to support more flexibility of teaching styles in conjunction with the user-modelling method are outlined. A complete list of CITE reports is included. (20 references) (Author/DB)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

X This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

E.S. Henderson

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

Guided discovery tutoring and bounded user modelling.

Mark Elsom-Cook,
Centre for Information Technology in Education,
Institute of Educational Technology,
Open University,
Walton Hall,
Milton Keynes,
MK7 6AA.

Abstract

The paper discusses a range of teaching styles possible for ICAI and suggests that they may usefully be thought of as points within a framework of guided discovery teaching methods. The use of Bounded User modelling within a guided discovery framework is introduced, and an application of the technique in a Lisp teaching system is described. Extensions to the system to support more flexibility of teaching styles in conjunction with this user-modelling method are outlined.

Appears in: Artificial Intelligence and Human learning J.Self (ed)
Chapman and Hall, 1988

Guided discovery tutoring and bounded user modelling.

1 Introduction

When we think of teaching, the first thought tends to be of an individual standing in front of a group of students holding forth on a particular topic. While acknowledging the existence of other approaches to teaching, practical considerations prevent us from affording them a central role in our current view of possible educational interactions. In particular, since our educational system has become oriented towards group teaching, we forget about the range of educational techniques which are possible in a one-to-one interaction between pupil and teacher. Socratic tutoring, free discovery learning, guided discovery learning, chalk and talk methods and programmed learning techniques are all examples of well-defined educational styles which can be (and have been) used in a one-to-one teaching situation.

One of the advantages of computer-based approaches to education is the possibility of providing a one-to-one correspondence between pupils and machines, facilitating individualization in learning. If we wish an ICAI system to offer a valuable educational experience on an individualized basis, we must require it to adapt the teaching style which it chooses to the needs of the pupil and the domain. There is a need for the system to reason about the style of tutoring which it performs, and to continually reassess the appropriateness of that style to the situation it is currently in.

A key dimension along which to categorise these teaching styles is the degree of constraint which is imposed upon the pupil. The range of actions which the pupil can perform are constrained by the environment in which she is operating. There are also constraints imposed by the tutor upon the pupil. It is the latter constraints being considered here. Such constraints may take a variety of forms, and the nature of the constraints imposed is determined by the style of tutoring which is adopted.

A free discovery interaction does not impose any of these constraints on the pupil. Any intervention by the teacher, however, is interfering with the course which the pupil would have followed, and hence is imposing an external constraint on the pupil (even advice which is ignored affects the pupil). It is therefore important, as part of the decision to take any teaching action, to be aware of the constraint being imposed upon the pupil and to be prepared to justify that constraint. This justification is likely to be in terms of the difference between the educational value of the course of action which the pupil was expected to follow and the course which she actually follows due to the intervention.

Although a real teacher will vary her teaching style as mentioned above, no existing ICAI systems are capable of this behaviour. Each system embodies a single teaching style, which is commonly a formalization of the intuitions of the author about teaching, rather than being derived from any firm educational grounding. The style which is chosen is normally highly constrained, with the system controlling the interaction strongly. This approach is adopted for practical reasons, to avoid dealing with some problems which are currently very difficult to handle. A primary reason is that it simplifies problems of communication with the pupil by allowing the tutoring system to maintain strong expectations about the possible range of student input.

Many systems claim to use some form of "Socratic dialogue" - a question and answer sequence directed towards uncovering the underlying misconceptions of the pupil. Although these systems have the surface form of such a dialogue, it is by no means always the case that the dialogue is truly Socratic. To structure such a dialogue for the benefit of the pupil requires Socrates [Broudy 1969] or an equivalent intelligence in the machine! Instead, most systems achieve a surface form which looks like such a dialogue without the underlying reasoning, and hence without the structuring which gives the dialogue it's "Socratic" value (though that does not imply that such dialogues have no educational value).

The diagram (Figure 1) is intended to be suggestive of the location of some existing systems on a dimension of constraint imposed upon the pupil: one end of the scale corresponds to totally "free" learning, while the other represents a strictly controlled style. Clearly, there are a range of forms of constraint, and the systems should really be organised in a space of constraint, rather than along a single dimension. The approximate ordering afforded by this illustration serves to give an overall feel for the range of tutoring systems, however.

For example, Anderson's Lisp tutor [Anderson 1985] is based on a very specific model of how Lisp should be learnt. The tutor tracks the actions of the pupil, and as soon as the pupil deviates from this "correct" learning path the tutor interrupts and moves her back onto that path. At the free extreme we find computers being used as free learning Logo environments in the manner advocated by Papert [Papert 1980] (though this is not necessarily how Logo is actually used).

Clancey's system GUIDON [Clancey 1979] is based around an expert system problem-solver in the domain of medical diagnosis. The tutor generates a problem-solution with this expert system, and assists the student in solving the same problem. Some experimentation which is not directly on the solution path is permitted, but the system immediately uses such experiments to initiate a dialogue which will bring the student back onto the "correct" solution path.

Stevens system WHY [Stevens 1974] permits the user more scope for directing the interaction. It initiates discussions on topics, and by default controls the interaction. If the pupil fails to respond directly to a question by the system, but chooses to introduce a related topic instead, WHY will permit the diversion, but remember the topic it was working on and return to it when the opportunity arises.

WEST [Burton 1982] acts as a "coach". It observes the student carrying out a task, and tries to develop a model of the pupil's ability to perform that task. The system sometimes offers advice to the pupil, but it does not necessarily correct the pupil every time it realises that she has made an error.

IMPART [Elsom-Cook 1985a] is a coaching system which provides guidance for pupils during the early stages of learning to program in Lisp. It observes a pupil interacting with a Lisp system, and provides advice about the semantics of the language only when it believes

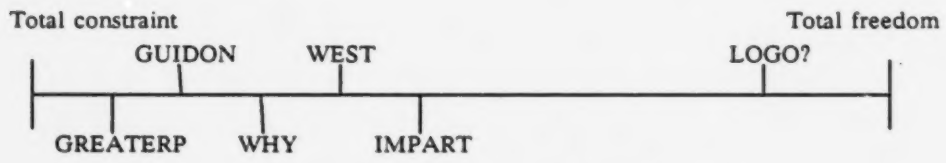

Figure 1: Constraint in teaching style

that the pupil has made an incorrect inference about the behaviour of a primitive Lisp function.

Optimally, we would like a tutoring system to be able to move backwards and forwards along this scale, using a variety of teaching styles which provide different levels and forms of constraint on the activities of the pupil. I wish to suggest that the whole range of tutoring styles may be thought of within the framework of guided discovery learning – offering differing degrees and forms of guidance involves choosing (and justifying) a point on this continuum of constraint which is appropriate to the current needs of the pupil and identifying teaching actions which provide the appropriate constraints. The sort of partially–directed tutoring (which has been called "coaching" [Burton 1982]) normally associated with guided discovery methods deals with a single point (or a small range) in this continuum of constraints. If we accept this redefinition of guided discovery, then moving between teaching styles corresponds to making a decision about the amount and form of guidance to provide.

The remainder of this paper will explore these issues in more depth by discussing the nature of guidance in tutoring and possible ways of providing such guidance (with particular reference to building models of the pupil). It concludes with a summary of the way in which such a model is used in a guided discovery Lisp tutoring system, IMPART [Elsom–Cook 1985a] and some suggestions for ways to achieve this flexibility of teaching style.

2 What is the nature of 'guidance' in tutoring?

R.S.Peters explains the essential idea behind guided discovery learning in the following manner:

> *" The cardinal function of the teacher, in the early stages, is to get the pupil on the inside of the form of thought or awareness with which he is concerned. At a later stage, when the pupil has built into his mind both the concepts and mode of exploration involved, the difference between teacher and taught is obviously only one of degree. For both are participating in the shared experience of exploring a common world. The teacher is simply more familiar with its contours and more skilled in finding and cutting pathways. The good teacher is a guide who helps others to dispense with his services."*
>
> [Peters 1966]

It is clearly not the case that pupil and teacher are always jointly exploring a domain as equals. Interactions which appear to be very similar on the surface may differ greatly when we explore the models which the participants in the interaction possess. Consider the following simple exchange:

A: " Can we talk about i/o now?"
B: " Sure, let's look at the 'write' predicate."

It may be the case that the tutor (B) has assessed the current knowledge state of the pupil with respect to the prerequisites and concepts associated with the domain, has decided that this area (i/o) is an appropriate one for further exploration, and has selected the 'write' predicate as a suitable topic for introducing the domain because it is not far beyond the

current state of knowledge of the pupil. On the other hand, B may simply have no clear idea about what to do next, and may have responded to the suggestion of looking at i/o by mentioning the only thing she knew about that domain.

Some child-centred approaches to education have run into difficulty by making the assumption that pupil and teacher are fellow-explorers, with no significant difference between them. The educational approach of Rousseau [Rousseau 1974], on the other hand, while retaining the basic idea of joint exploration/guidance, represents the opposite extreme in that it invariably assumes that the teacher has "the correct version" of what is to be learnt, and constantly restructures the learning experience of the pupil in order to guide her to that model (which corresponds to enforcing a goal of the teacher on the pupil). Such an approach is clearly still child-centred in that the focus is firmly on adapting the educational experience to the individual (since the teacher's goals are determined by the state of the pupil), but it acknowledges an "ultimate authority" vested in the tutor on the grounds of her greater knowledge of the domain. It is not the case that there will always be a "correct version" of the thing which the pupil is learning, or that, even if such a correct form exists, the teacher will know it. As the pupil becomes more proficient in the domain the authority of the tutor will diminish, and it is important that this should be reflected in the interaction. Our tutor must, therefore, be able to act appropriately both as guide and as fellow-explorer with the pupil.

A guided discovery approach to learning a new domain, then, will take the pupil along the continuum of constraint from a heavily structured, tutor-directed interaction to one in which the tutor plays no role. How this transition occurs depends upon specific details of the pupil, domain and teacher. To explore this further it is necessary to be more explicit about what guided discovery learning methods mean. In particular, what is the permissible range of guidance methods which the system should embrace, and how can it decide about making a transition from one method to another? At the free extreme of the scale, the tutor will be completely invisible, leaving the student to explore the task being taught. Clearly, the tutor is still present, and must be observing the actions of the pupil in order to decide whether changing to an approach with more intervention is necessary. At the other extreme, the tutor controls the interaction completely, but should be eliciting information from the pupil which is used to adjust the interaction to match the abilities of the pupil. Using these styles as the end-points for our teaching methods explicitly excludes "free" learning at one end of the scale, and rote learning methods/pure presentation methods etc. at the other end. These can reasonably be excluded on the grounds that they are not teaching interactions. The former because no teacher is present, and the latter because no account is taken of the presence of the pupil; though this is not to deny that they may be valid educational experiences.

Within this framework we can regard the teaching interaction as a negotiation between two agents, one of which (the tutor) brings domain-specific knowledge and teaching techniques to the interaction, while the other (the pupil) brings various pre-existing cognitive structures and preferred learning methods (Figure 2). If we are considering a situation in which the pupil wants to learn, and the teacher wants to teach, then these agents have the joint goal of imparting the domain-specific knowledge which resides in the teacher's head to the pupil. This requires changing that knowledge from the form in which it is useful to the teacher (i.e. as applied expertise in the domain) into a form suitable for communicating, transferring that knowledge to the pupil, and ensuring that the knowledge is fully integrated with the existing cognitive structures of that pupil. Any inconsistencies between the pre-existing knowledge of the pupil and the new information must be resolved. It must also be ensured that the

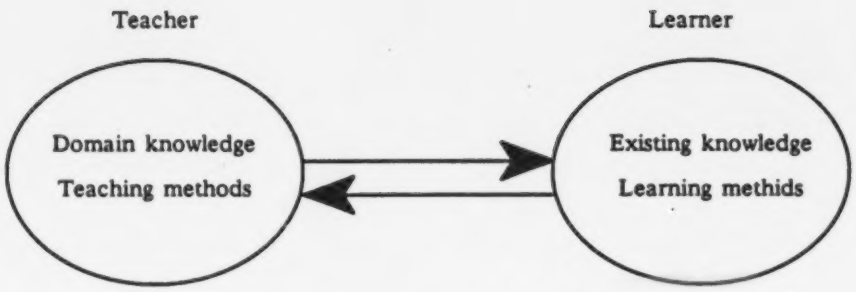

Figure 2: Teaching interaction

knowledge is integrated in such a way that it can be applied appropriately, rather than simply being "inert" information [Whitehead 1932]. This task of integration is often more important and difficult than the task of transfer. In achieving the form of interaction discussed above, the teacher restructures the domain expertise using knowledge about the teaching process, the learning processes of the pupil, the current state of the pupil and, optionally, some domain-specific expertise about the common difficulties of students in this domain. Few of these forms of knowledge are represented in existing tutoring systems. It should be clear from this discussion that an expert system approach to ICAI is inappropriate. Knowledge is encoded in such systems in order to be used, rather than in order to be imparted to others [Clancey 1986] and achieving an appropriate restructuring for tutoring (which is dependent on the particular pupil being taught) requires a deeper representation of the domain than that which expert systems technology can provide.

As described above, Guided discovery teaching requires the system to constantly reassess the appropriateness of the teaching activities which it is performing, monitoring the success (and possible future choices) of the pupil, and (either by direct or indirect methods) ensuring that the pupil follows paths corresponding to a fruitful educational experience (avoiding blind alleys etc).

In achieving this, I wish to suggest that the teacher has four major forms of goal which must be satisfied by the guidance being given: presenting new material, finding out about the current state of knowledge of the pupil, maintaining an appropriate form for the teaching interaction itself, and maintaining/developing the motivation of the student.

These goals can be expanded into more detailed goal structures which are dependent on the current state (and history) of the tutoring interaction (and the pupil). The tutoring task then becomes one of (partially) satisfying multiple goals simultaneously. One utterance may provide input to many of these goals. This task is closely related to the sort of models of goal-directed dialogue which have been developed recently (e.g. [Perrault 1980] [Appelt 1982]) and it seems natural to regard the tutoring interaction as being a special application of more general dialogue mechanisms.

Mark Miller [Miller 1982] offers a flavour of the issues which we need to explore in deciding on appropriate guidance by presenting the following set of questions:

1) What constitutes a good hint?
2) Should the computer spontaneously interrupt the student to provide remedial advice?
3) How often should the tutor intervene?
4) How much should it say?
5) When should it say it?
6) How should it say it?

Most of these questions are associated with the goal of presenting material appropriately. This reflects the fact that tutoring systems have tended to focus on this particular aspect of the tutoring problem.

3 How can appropriate guidance be offered?

Providing appropriate guidance for a pupil is difficult. There are many unresolved educational issues in this area, of which the most important is probably that of control in the interaction. Who controls the interaction has been a topic of discussion in education theory for many years, but since the move away from one-to-one education it has become less central. ICAI systems bring this issue to the foreground once again. Under what circumstances is it justifiable for the tutor (human or computer) to take control of the interaction and direct the pupil - possibly in directions contrary to those which the pupil wishes to pursue. In terms of the two agent model described above, this becomes an issue of resolving goal conflict between the agents. The teacher cannot control the interaction without gaining the cooperation of the pupil in achieving that goal - either by providing a justification for following a particular goal in terms the pupil can understand (which Whitehead [Whitehead 1932] claims is a basic prerequisite of any educational decision) or by appealing to some higher level goal in the interaction (e.g. "You said you wanted to learn this, so trust me for half-an-hour and you will!").

Of central importance to the provision of this negotiated tutoring interaction is an understanding of the problems which confront the pupil at a particular time. Offering appropriate information in forms comprehensible to the pupil at the right times, generating suitable problems, providing useful feedback on progress, maintaining motivation etc. all require some model of the current state of knowledge (and emotion!) of the pupil. ICAI work has focussed on building models of the pupil to permit selection of appropriate teaching material, but another important aspect of the user-model has often been neglected. I wish to suggest that deciding what to teach next must be based NOT primarily on what the teacher perceives as the logical structure of the domain, but on the options which the pupil perceives to be within her capabilities at a particular time. This involves maintaining some model of the reasoning processes which the pupil uses to acquire knowledge. These processes include inductive reasoning, analogical reasoning, problem-solving, theory-testing etc. There will be both domain-independent and domain-dependent learning skills involved.

3.1 Providing and using models of the learning process.

Most existing ICAI systems build models of the pupil as variants of an "expert proficiency" in the domain. There is assumed to be some expert competence which can be broken down

into a set of skills so small that the pupil either has them or doesn't. The task of the tutor is then to fill the student with these skills like filling a bucket with knowledge. This can be described as "subset modelling" (figure 3) (it is often known as "overlay modelling"). The alternative expert-based approach is to include a second bucket ("bugs") containing tiny subskills which the expert doesn't have, which produce incorrect behaviour, but which have been observed to exist in pupils learning about the domain. The tutor must now pour in subskills from the first bucket while removing any subskills which the pupil has acquired from the second bucket. This may be described as "perturbation modelling" (figure 3). These approaches both imply a very simplistic model of the learning process (not far removed from rote learning), which takes no account of the rich range of learning styles and capabilities for which there is psychological evidence.

If we move away from the common expert-based modelling methods then we should replace them with an approach in which the model is constructed in terms familiar to the learner, with reference to the range of reasoning processes mentioned above. I wish to propose that an appropriate framework for this activity is that of bounded user-modelling (see figure 4). In this type of model, the understanding of the state of the pupil exists as a set of upper and lower bounds on the possible states of the learner. The tutor will not necessarily know

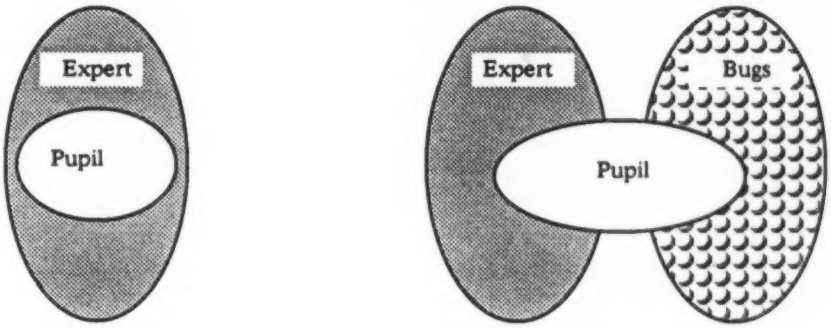

Figure 3: Expert-based modelling methods

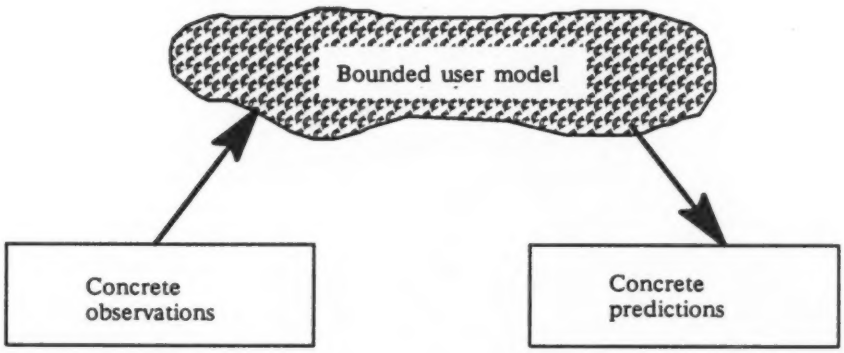

Figure 4: Bounded user modelling

"exactly" what state the pupil is in. Such "exact" knowledge is not generally possible (certainly human teachers do not use it), and much teaching can be done using approximations to the state of the pupil.

The bounds on the model are generated by applying a model of the reasoning activities which the student uses when confronted with some event in the world. Obviously the model of reasoning processes in the system will not match that of the student perfectly, and avoiding major errors involves the system in a cycle of generating bounds from concrete observations which the pupil could make, then testing them by generating some concrete predictions from those bounds and comparing them with what actually happens (an example is given below). The student's reasoning processes are themselves changing over the course of time (though not with the frequency of the domain model) and should themselves be the subject of user-modelling activities (these changes reflect the student learning about learning in this domain). It may be the case that there are more levels of change (learning to learn about learning etc.), but since each level is considerably more stable than the last, we will not consider these higher levels any further.

The use of upper and lower bounds in the Bounded user model is related to Mitchell's version space (a way of formally representing the range of permissible descriptions of an event in terms of a greatest and least member of the set of descriptions) [Mitchell 1981] but is different in several respects. In Mitchell's approach, a focussing algorithm is used which generates the most general and most specific descriptions of an event. Subsequent events are used to focus on the correct concept being learnt, by making the general description more specific and the specific description more general. These two descriptions will converge as more events are observed. In IMPART, there may be a number of things to be learnt from each event, and consequently each event may generate a completely new upper and lower bound. There is not a one-to-one mapping between events and concepts to be acquired. In effect the system is maintaining a set of version spaces. This approach has the advantage of permitting the use of multiple partial models of the learning process, and allows the representation of inconsistent beliefs held by the pupil (though these advantages are not explored in the implementation discussed below). The price paid for this approach is the loss of the partial ordering of descriptions on which Mitchell's focussing algorithm relies. This results in an increase in the complexity of the algorithm which must be used to obtain descriptions of events, and a greater range of possible models to examine. How IMPART deals with this problem is further discussed below.

Moving towards learner-based modelling methods was also a motivation behind the development of the Genetic Graph [Goldstein 1982]. This data-structure represents a domain by linking the domain knowledge with arcs describing the form of learning which must take place in order to move from one piece of knowledge to another. As originally described this approach involved making these links static and explicit (in effect embodying a theory about one possible way to learn the domain). An individual student is represented as a subset of the genetic Graph. In many ways such an approach is like the expert-based modelling methods described above - the difference being that the pupil is now being compared with one particular expert learner instead of one particular domain expert. It does not provide any scope for extension such that individual differences in learning style, or developmental changes in the pupil (other than with respect to domain knowledge) can be handled. The original version was intended for modelling acquisition of procedural rules, and though Goldstein has subsequently discussed extensions to this approach to support

multiple representations, declarative knowledge, planning etc, this work has not been pursued.

As has been observed, creating and maintaining a Bounded User model relies on having some model of the reasoning processes which the pupil is applying in order to organise the environment. This dependence on models of inductive and deductive reasoning processes appears to cause a problem for this approach. At present, there are no suitably complete and formal models of these processes available. Psychological theories in this area are normally incomplete, rarely formalized and (except in special cases) can only be shown to have very restricted applicability.

If the task which we are attempting to perform with our tutor is a pure learning task, as would be the case at the "joint exploration" end of the path of development discussed earlier, then this shortage of theories severely restricts the usefulness of the tutor. The tutor cannot guide the pupil in a manner that avoids dead-ends and unproductive courses of study, because the tutor cannot assess such problems, and will get drawn into the same traps as the pupil. The tutor cannot serve usefully as a co-explorer because this shortage of theories will render it less effective than the pupil, and may result in it leading the pupil away from an appropriate course of exploration due to its lack of knowledge. For earlier parts of the tutoring process, however, the problem is not so serious because the tutoring system can be provided with some description of what it is that the pupil is expected to learn. Since the tutoring system has such knowledge, it has some "target" against which to compare the hypotheses generated by the model(s) of the learning process, and some way to prevent those hypotheses becoming too far divorced from useful models of the domain. In consequence, less reliance is placed on the completeness or consistency of the learning models which are used. This approach places IMPART much closer to Rousseau in terms of the approaches it can support than more naive pupil-centred approaches.

The availability of a benchmark model describing what the pupil is expected to learn facilitates the use of multiple partial models of the learning process rather than requiring a single correct model, because any discrepancy between the predictions of different learning models can be resolved by reference to this model. This avoids the "joint exploration" problem of permitting tutor and pupil to wander hand in hand up blind alleys and become hopelessly stuck. There seems to be no clear evidence for the existence of a single learning mechanism in any case, so permitting the tutor to "mix and match" among multiple learning models is likely to permit a closer approach to an accurate model of the pupil.

The best formalized models of learning which currently exist are those derived from the machine learning literature. These models are currently quite simple, applicable only to special domains, and (generally) have no particular claim to relate to human learning. Any attempt to use these systems directly to model a learner would fail for these reasons, but used at one remove to create possible bounds for the model of the internal state of the student, such machine learning approaches are useful. The following section discusses the way in which such techniques are applied in an existing experimental tutoring system.

3.2 Machine learning in IMPART.

IMPART [Elsom-Cook 1984] is a tutoring system which applies guided discovery methods in the domain of programming. It teaches the student about the syntax and semantics of

Lisp through experiments with a Lisp interpreter. It assumes that the student initially knows nothing about the language, and applies the bounded user modelling technique described above to develop a model of the pupil which contributes to decisions about the actions which the tutor should take. This bounded user-model is one of the primary inputs to a goal-directed dialogue system which structures the educational interaction (for a sample interaction see [Elsom-Cook 1985a]). The system only supports the early stages of learning Lisp, and has not been used with real pupils. It does not achieve the wide variation of teaching style that was discussed earlier, but supports a small range of guidance methods.

The learning model which is used is a simple inductive learning algorithm based upon standard machine learning methods [Michalski 1983] No attempt to change the learning mechanisms over time is made.

IMPART is provided in advance with a set of descriptors that can be used to describe the possible attributes of an event in Lisp, and a set of descriptions of the actual behaviour of each Lisp function couched in terms of those descriptors. When an event occurs (such as the evaluation of a lisp expression), the tutor uses the descriptors to produce a description of the observable behaviour corresponding to that event. This is the lower bound on the student model.

Having established this description of the surface form of the event, the system then does a depth-restricted exhaustive search through a set of inductive inference and transformation rules in order to derive the most general description of the event which is consistent with the learning model which it uses. IMPART embodies a subset of the algorithm mentioned above in a set of 7 inference rules which can be applied at any time. It also includes a set of transformation rules which may be regarded as corresponding to domain-specific learning methods and heuristics for making sense of the world. They consist of suggestions for discovering relationships between things (for example, if some Sexpression whose origin is unclear appears, then evaluate all the available atoms to see if it the value of one of them). The details of these rules can be found in [Elsom-Cook 1985c].

These two descriptions (surface and general) are actually a maximally-specific conjunctive description and maximally-general conjunctive description of the event, and can be regarded as lower and upper bounds on what the student would infer from this example. Regarding these bounds in this way is obviously an assumption about the nature of the learning processes which the pupil is using; it is quite possible to imagine a range of idiosyncratic inferences which the student could make and which do not fit into this model, but IMPART does not approach such issues.

The pair of bounds which are generated are compared with bounds generated by other events. If the new bounds are more specific versions of older ones, then the older ones are removed. If they are more general then the newer ones are removed. Since the partial-ordering of the version space has been lost, not all bounds can be assigned to one of these categories, and in this case IMPART simply leaves the bounds in the model. This is one area where the system requires much further work.

Of course, the inferences made by this system are normally more general than those which the pupil makes, and are often wrong (since induction is not truth-preserving). The problem of making the correct inductive inferences and narrowing the bounds is difficult, and IMPART uses two mechanisms to avoid dealing with the full learning problem.

Firstly, the inference process is constrained in two ways; by limiting the depth of inference (i.e. number of rules which can be applied), and by using a number of domain-specific

heuristics (which can be taught to the pupil) that restrict the search space by making use of knowledge about the likely range of descriptions of Lisp events.

The second mechanism is the deductive reasoning model shown in Fig. 4. This is a simple forward-chaining deductive mechanism which is applied to the descriptions in the Bounded user-model to generate a range of predictions about the behaviour of other expressions. As the model changes, the predictions change. These predictions are used to generate problems for the pupil to solve, together with predicted answers whose origin can be traced back to a specific part of a specific description in the model. If the pupil solves the problem, the solution can be compared with the predictions and used to update the model. A sample of problem-generation appears in [Elsom-Cook 1985b].

To illustrate this approach, consider the model which would be generated from a very simple Lisp event, assuming no previous Lisp experience of the pupil:

```
EVENT:  (QUOTE IRENE)
        IRENE

Lower bound:
  quote(irene)
    [type(quote(irene),sexpr),
     type(irene,atom),
     out(value,irene)]

Upper bound:
  quote(ARG1)
    [type(quote(ARG1),sexpr),
     type(ARG1,sexpr),
     out(value,ARG1)]
```

In this case, the lower bound is purely a description of the surface form of the event (i.e. things that are directly observable). The information about types of objects, and about function/argument distinction which are used in this description are things which would be flagged as being necessary to teach directly to a naive pupil. The upper bound is (coincidentally) a correct description of the behaviour of QUOTE. Between these descriptions lie many possible errorful descriptions such as QUOTE being a function with one argument which always returns the value IRENE, or which can only be applied to atoms. If the system were to ask the student what behaviour the event (QUOTE (EAT FROGS)) would cause, the following predictions would be made:

```
Lower bound prediction: Dunno,nothing,error
Upper bound prediction: (EAT FROGS)
```

Examples of intermediate cases (which are not actually generated, but which can be derived from the bounds if the pupil makes a prediction other than those given above) are:

```
Returns value irene:  IRENE
Applies to atoms: That's not an atom, error
```

The learning model in IMPART is extremely naive and simplistic, even by the standards of machine learning theories let alone from the psychological viewpoint. However, it does produce useful behaviour and is a major factor contributing to the smoothness of interactions with the system. Further development of this approach should include supporting multiple models, permitting modelling of the learning processes of the student, and enriching the deductive reasoning processes used by the tutor.

3.3 Bounded user modelling and flexibility of teaching style.

The central idea behind bounded user modelling is that it is not possible to create a completely accurate model of what is inside the pupil's head, and that any attempt to do so cannot succeed. Human teachers do not appear to build models of this type, but they DO build models. These models provide approximations to the student's knowledge.

An approximation to the knowledge of the student provides a sound basis for teaching, since one particular teaching act is often applicable to a range of states within the pupil, and it is only necessary to ensure that the pupil is in one of those states in order to know that the teaching act is appropriate. There is a close link between the level of detail of the model and the amount of control which can be exercised in teaching style. Only very accurate models are adequate for highly directive teaching styles, but guidance can be given with less detailed knowledge.

Because the uncertainty about the knowledge of the pupil is explicitly represented in the Bounded User model, it provides a limit for the range of teaching styles which can be applied at a particular moment. Highly directed styles cannot be used if the knowledge about the pupil in that area is very vague, for example. IMPART only had one teaching style, comprising about 6 separate teaching actions, and so made little use of this facility. I wish to suggest that a wider variety of teaching actions, each associated with information about the level of detail needed in the model would provide the basis for a mechanism of selecting between a range of teaching activities as was discussed earlier. IMPART did not contribute to such a mechanism, but this is an area in which further research is being conducted.

4 Conclusions.

In summary, I wish to suggest that tutoring styles should be thought of as existing as points along a continuum of guided discovery teaching styles. Tutoring systems should aim to cover a range of adjacent styles on that continuum, and be prepared to switch between them. It is important that guided discovery learning systems possess domain-specific knowledge which can be used to ensure that the pupil does not move too far away from fruitful models of the domain. The user-modelling in these systems should relate to models of the learning process, and should operate by imposing fuzzy bounds on the knowledge of the pupil rather than attempting to build an exact model. Given the current state of machine learning techniques, some simple versions of such an approach can be implemented, but a richer range of formal models of learning is ultimately required.

References

[Anderson 1985]
J.R.Anderson,B.J.Reiser - The Lisp tutor.
BYTE April 1985

[Appelt 1982]
D.E.Appelt – Planning natural language utterances to satisfy multiple goals.
Tech note 259 SRI International 1982

[Broudy 1969]
H.Broudy – Socrates and the teaching machine.
in The Human Encounter
Eds. S.Stoff and S.Schwartzberg
Harper–Row 1969

[Burton 1982]
R.R.Burton,J.S.Brown – An investigation of computer coaching for informal learning activities.
in Intelligent Tutoring Systems
Eds. D.Sleeman and J.S.Brown
Academic Press 1982

[Clancey 1979]
W.J. Clancey – Dialogue management for rule–based tutorials.
Proceedings 6th IJCAI 1979

[Clancey 1986]
W.J.Clancey, R.Letsinger – Neomycin: Reconfiguring a rule–based expert system for application to teaching.
in Expert systems in medicine
W.J.Clancey, E.H.Shortliffe (eds)
Addison–Wesley 1986

[Elsom–Cook 1984]
M.T.Elsom–Cook – Design considerations of an Intelligent tutoring system for programming languages.
Ph.D thesis. University of Warwick.
(CITE Report 8 Open University) 1984

[Elsom–Cook 1985a]
M.T.Elsom–Cook – Design considerations of an Intelligent tutoring system for programming languages.
in Human Computer Interaction: Interact '84
B.Schackel (ed)
Elsevier 1985

[Elsom–Cook 1985b]
M.T.Elsom–Cook – Towards a framework for human–computer dialogue.
in People and computers: Designing the interface.
P.Johnson, S.Cook (eds)
Cambridge University Press 1985

[Elsom–Cook 1985c]
M.T.Elsom–Cook – Machine learning and user modelling.
CAL research group report. Open University. 1985

[Goldstein 1982]
I.P.Goldstein – The genetic graph: a representation for the evolution of procedural knowledge.
in Intelligent Tutoring Systems
Eds. D.Sleeman and J.S.Brown
Academic Press 1982

[Michalski 1983]
R.S.Michalski – A theory and methodology of inductive learning.
in Machine Learning.
eds.R.S.Michalski,J.G.Carbonell,T.M.Mitchell.
Tioga 1983

[Miller 1982]
M.L.Miller – A structured planning and debugging environment for elementary programming.
in Intelligent tutoring systems.
eds. D.Sleeman,J.S.Brown.
Academic 1982

[Mitchell 1981]
T.Mitchell – Generalization as Search
Artificial Intelligence 1981

[Papert 1980]
S.Papert – Mindstorms.
Harvester Press 1980

[Perrault 1980]
C.R.Perrault,J.F.Allen – Analysing intention in utterances.
Artificial Intelligence 15 1980

[Peters 1966]
R.S.Peters – Ethics and education.
Allen and Unwin 1966

[Rousseau 1974]
J.J.Rousseau – Emile. (1762)
Everyman edition. 1974

[Stevens 1982]
A.Stevens,A.Collins,S.E.Goldin – Misconceptions in students understanding.
in Intelligent Tutoring Systems
Eds. D.Sleeman and J.S.Brown Academic Press 1982

[Whitehead 1932]
A.N.Whitehead – The aims of education.
Benn 1950

CENTRE FOR INFORMATION TECHNLOGY IN EDUCATION

List of CITE Reports

These reports may be obtained from:

Hansa Solanki, Institute of Educational Technology, The Open University, Walton Hall, MILTON KEYNES, MK7 6AA, England.

Report No.	Title and Author
1	A.T. Vincent, (1985) Computing and the Blind.
2	A. Jones, G. Kirkup, J. Morrison (1985) A Trial of Home Based Computer Terminals.
3	Gill Kirkup, (1985) The Present and Potential Use of Ceefax in the Open University.
4	Mark Elsom-Cook, (1986) Artificial Intelligence and Computer Assisted Instruction.
5	Mark Elsom-Cook, (1986) A Pascal program checker.
6	Simon Holland, (1986) How computers are used in the teaching of music and speculations about how Artificial Intelligence could be applied to radically improve the learning of composition skills.
7	Simon Holland, (1986) Design consideration for a human-computer interface using 12-tone three-dimensional harmony space to aid novices to learn aspects of harmony and composition.
8	Alison Petrie-Brown, (1987) The Influence of Context and Coherence as a Foundation for Dialogue Research.
9	Eileen Scanlon, Randall B. Smith (1987) A Rational Reconstruction of a Bubble Chamber Simulation Using The Alternate Reality Kit.
10	Mark Elsom-Cook, (1987) Intelligent Computer-Aided Instruction research at the Open University.
11	Mark Elsom-Cook, (1987) Towards a framework for human-computer discourse.
12	Mark Elsom-Cook, (1987) MATILDA AND IMPART: Lisp tools.
13	Mark Elsom-Cook, (1987) Guided discovery tutoring and bounded user modelling in Intelligent Computer Aided Instruction ed. J. Self, Chapman-Hall 1987.